A Probabilistic Approach to Trajectory Generation in the Presence of Uncertainty

Jeffrey B. Jewell
NASA Jet Propulsion Laboratory
Pasadena, CA 91109
Email: Jeffrey.Jewell@jpl.nasa.gov

Seungwon Lee
NASA Jet Propulsion Laboratory
Pasadena, CA 91109
Email: Seungwon.Lee@jpl.nasa.gov

Martin Lo
NASA Jet Propulsion Laboratory
Pasadena, CA 91109
Email: Martin.Lo@jpl.nasa.gov

***Abstract*—Achieving longer mission duration, increased science return, and flexability for follow on observations motivated by unanticipated discoveries is facilitated through improvements in trajectory generation, particularly with respect to computational efficiency of computing near-optimal solutions of constrained optimization problems. Practical complications to be faced also include solving an ensemble of constrainted optimization problems for random initial and/or final states, uncertainty in the dynamics, and quantifying state measurement error during tracking and navigation. Here we formulate a probabilistic approach to control allowing convergence to optimal solutions over an entire set of boundary conditions, as well as allowing other sources of uncertainty to be included. Furthermore, we use insight provided from a global view of phase space structure (i.e. a "dynamical systems" viewpoint of the free trajectories) to guide computation and improve efficiency of trajectory generation. A numerical example of providing initial guess solutions for the restricted three-body problem is given, and future work including stochastic algorithms for improvement of these initial solutions discussed.**

I. Introduction

Increasing science return from space missions, through longer mission duration and flexability for follow on observations motivated by unanticipated discoveries, is facilitated through improvements in trajectory generation, particularly with respect to computational efficiency of computing near-optimal solutions. Practical complications to be faced include solving an ensemble of constrained optimization problems for random initial and/or final states, including uncertainty in the dynamics, and quantifying state measurement error during tracking and navigation. Here we formulate a probabilistic approach to control allowing convergence to optimal solutions over an entire set of boundary conditions, as well as allowing other sources of uncertainty to be included. Such a capability is particularly important in both the design stage (where typically a "sensitivity analysis" is important in order to validate the capability of the spacecraft to achieve the mission objectives even in the presence of launch vehicle error and other perturbations) and in flight, requiring precise control over uncertainty for tracking and navigation.

Ideally, mission design could begin with scientists listing a collection of "interesting way points" (a rank ordered prioritized list of scientifically interesting targets, with times on station, etc.). The problem then is to solve for an *optimal trajectory*, which hits as many of the way points of interest with the smallest expenditure of, for example, fuel. The importance of finding optimal solutions to the trajectory design problem is an increased capacity of potential science return (quantified intuitively by number of "interesting places" visited, and length of time "on station" at each). Given constraints such as finite time and fuel, there is obviously a maximum value of the "potential science return". Feasible solutions possibly far away from globally optimal solutions, may be entirely acceptable for some less ambitious missions with a small number of objectives. However, for missions the ambitious goal of maximizing potential science return, it is increasingly important to have algorithms which can explicitly control global distance to optimal solutions.

Optimal control problems are notoriously difficult to solve for nonlinear dynamical constraints (see for example Bryson and Ho, 1969). However, what has recently been recognized and applied (Howell et al. 1997, Serban et al. 2002, Gómez et al., 2004) to current and future mission design (including Genesis, WMAP, and Planck scheduled for launch in 2008) is the ability of dynamical systems theory to provide insight and guide computation of optimal trajectories. Specifically, the science goals of these missions have been enabled by using halo orbits about the Lagrange point L1 of the Earth-Sun-spacecraft three-body dynamics. These halo orbit of the Genesis mission was the first to be designed entirely from dynamical systems theoretic insight This work validated the concept of trajectories which follow invariant manifolds of the underlying dynamics, and led to the general idea of the "interplanetary superhighway", in which trajectories can be designed by hopping on and off the manifolds of the dynamics providing free or extremely low cost transport throughout 3 or higher body systems. A wide range of mission concepts have subsequently been proposed which make use of these trajectories, and re-revitalized the goal of computing optimal or near optimal trajectories for space mission design. Provided we can meet the challenges of computing these trajectories and subsequently navigating them, *we are rewarded with longer mission duration and increased potential for science return*.

Compounding the challenge are sources of uncertainty that require not just high accuracy and "near-optimal" solutions for one set of initial and final target states, but an entire ensemble of these problems due to uncertainty in initial conditions (such as arising from launch vehicle error). The importance of characterizing the entire set of optimal solutions for random boundary conditions is important for validation of a mission design to achieving a majority of the objectives even in the presence of various uncertainties, including launch vehicle error and/or perturbations to the model of the dynamics. Halo orbits for example, have the unfortunate characteristic that their successful navigation is extremely sensitive to launch vehicle error . In contrast to interplanetary missions where launch vehicle error can be corrected within 7-14 days after launch, halo orbit missions must generally correct launch vehicle error within the first 7 days after launch or the required ΔV to correct the trajectory will be beyond the spacecraft's capability (Serban et al, 2002). We therefore need to control the computational error over an entire ensemble of boundary conditions in order to accommodate initial periods of spacecraft checkout and orbit determination which invariably cut into our overall fuel budget.

Precise quantification of uncertainty in the actual trajectory is therefore increasingly important in order to accurately compute and closely navigate optimal trajectories utilizing the underlying "connectivity of phase space" provided by the full underlying nonlinear dynamics. For this purpose we therefore need an algorithmic framework which provably converges to optimal solutions for randomly chosen initial and/or final states. It is difficult to prove convergence of deterministic algorithms for optimal control problems with nonlinear dynamical constraints due to the potential convergence to local as opposed to global minima of the objective function defining optimality. *The goal of convergence on sets of optimal solutions with randomly chosen boundary conditions leads us to consider a probabilistic approach to the entire problem itself.* Moreover, other sources of uncertainty can naturally be included in this framework, including perturbations to the dynamics, as well as the closed-loop control challenge of smoothing (quantifying uncertainty in the entire past trajectory, up to and including the current state) when supplied with state measurements with noise.

Our strategy is to 1) use insight provided by dynamical systems theory to quickly sample from what we call the "initial ensemble" of solutions, 2) construct a sequence of probability measures directly which provably converge on sets of optimal solutions given randomly chosen boundary conditions, and 3) use sampling algorithms such as Markov Chain Monte Carlo or particle filters to directly sample solutions from these probability densities, with increasingly better solutions "learned" as computation progresses.

While more expensive than standard deterministic schemes for single BC's, as discussed above, *the goal is really to control the figure of merit defining optimality over entire sets of solutions, in order to permit validation (during design) and implementation (during flight) of mission scenarios in the presence of uncertainty*. As mentioned above, the potential payoff of rendering the probabilistic approach computationally efficient enough for practical use are trajectories which offer the capacity for increased science return through increased mission duration and added flexibility for follow-on observations motivated by unanticipated discoveries.

In this paper, we first review the formulation of trajectory design as a problem of optimal control, and discuss the resulting algorithmic challenges in solving these problems. We then discuss the three-body dynamics and the dynamical systems theoretic insight leading to near-optimal trajectories. We then provide an overview of the probabilistic framework for these problems, and comment on progress made in the first step of initializing the samples of solutions for sets of random boundary conditions. We close with a discussion of work in progress leading to stochastic algorithms for improvement of these trajectories and convergence in probability to sets of optimal solutions for random boundary conditions.

II. Overview of Trajectory Design and Navigation

A. *Statement of the Problem*

Here we provide a brief overview of the trajectory design problem. We are interested in finding control inputs $u(t)$ such that for some dynamics $F \circ y$ (with y the system state vector and $F \circ y$ the vector field of the dynamics evaluated at the current state) the trajectory, given as a solution to the non-autonomous ordinary differential equation (ODE)

$$\dot{y} \quad = \quad F \circ y(t) + u(t) \tag{1}$$

traverses from some specified initial state $y(0)$ to some specified target final state $y(T)$.

As an example for this paper, we consider the dynamics of the (circular restricted) three-body problem including control inputs involve the system of equations (see for example Serban et al. 2002)

$$\begin{aligned}
\dot{y}_1 &= y_4 && (2)\\
\dot{y}_2 &= y_5 && (3)\\
\dot{y}_3 &= y_6 && (4)\\
\dot{y}_4 &= 2y_2 + \frac{\partial U}{\partial y_1} + u_1(t) && (5)\\
\dot{y}_5 &= -2y_1 + \frac{\partial U}{\partial y_2} + u_2(t) && (6)\\
\dot{y}_6 &= \frac{\partial U}{\partial y_3} + u_3(t) && (7)\\
m(t)u_1(t) &= \dot{m}(t)V_{nozzle}\sin\theta(t)\cos\phi(t) && (8)\\
m(t)u_2(t) &= \dot{m}(t)V_{nozzle}\sin\theta(t)\sin\phi(t) && (9)\\
m(t)u_3(t) &= \dot{m}(t)V_{nozzle}\cos\theta(t) && (10)
\end{aligned}$$

where in the rotating coordinate system

$$\begin{aligned}
U &= \frac{1}{2}(y_1^2 + y_2^2) + \frac{1-\mu}{[(y_1+\mu)^2 + y_2^2 + y_3^2]^{1/2}}\\
&\quad + \frac{\mu}{[(y_1 - 1 - \mu)^2 + y_2^2 + y_3^2]^{1/2}} && (11)
\end{aligned}$$

and where in the above we have assumed the control inputs are generated by a rocket with nozzle velocity V_{nozzle} and which can swivel through angles (θ, ϕ), and finally can be throttled by controlling $\dot{m}$.

The true solutions of the dynamics, otherwise known as the *free trajectories*, are simply those which, for a given initial condition $y(0)$ have $u(t) = 0$. Clearly, the set of potential target states $y(T)$ is limited by the restriction of vanishing control input. Provided we have a "powerful rocket", we can effectively go anywhere in the phase space desired, and in fact, for any chosen path, the control input required to follow that path is trivially computed by re-arranging the above to *solve* for $u(t)$ according to

$$u(t) \quad = \quad \dot{y} - F \circ y(t) \tag{12}$$

Note that in what follows we will refer to $u(t)$ either as the "residual", or "defect", or "control input" ($u(t)$ is one measure of error typically controlled when numerically solving initial value problems for ODE's). Typically, the *cost* of following any given path can be related to the "magnitude" of the control input, measured with a norm on the space of continuous functions

$$\|u\|_{L^q} = \left(\int_0^T dt \ \|u(t)\|^q_{l^p} \right)^{1/q} \tag{13}$$

with $\| \cdot \|_{l^p}$ a *vector* norm (i.e. the norm of the instantaneous control vector at the time t).

Denoting the cost of traversing the path $y(t)$ from specified boundary conditions $y(0) = x_i$, $y(T) = x_f$ in time T

$$\Gamma(y(x_i, x_f)) = \|\dot{y}(t; x_i, x_f) - F \circ y(t; x_i, x_f)\|_{L^q} \tag{14}$$

our goal is to find an optimal path y^* so that

$$y^*(x_i, x_f, T) = \min_y \Gamma[y(x_i, x_f, T)] \tag{15}$$

In order to compute such an optimal path we can turn to a formulation of optimal open-loop control using the calculus of variations, which we now briefly discuss (we will use this later in a simplified linear setting later in order to efficiently compute initial approximations to the y^* which can then be improved with subsequent computation).

B. Optimal Open-Loop Control

Following (Bryson and Ho, 1969), optimal (open-loop) control is typically formulated as constrained optimization, with the associated "action" functional

$$\begin{aligned} J \quad = \quad & \phi[x_f, T] + \nu \cdot y(T) + \sum_i \Big[\lambda \cdot y|_{t_i}^{t_{i+1}} \\ & + \int_{t_i}^{t_{i+1}} dt \ L(y, u, t) + \lambda(t) \cdot (f(y, u, t) - \dot{y}) - \dot{\lambda} \cdot y \Big] \end{aligned}$$

(where in the above $f(y, u, t) \equiv F \circ y + u$) which gives, from the calculus of variations, the equations

$$\begin{aligned} \dot{y} &= f(y, u, t) \\ \dot{\lambda} &= -\frac{\partial L}{\partial y} - \frac{\partial f}{\partial y} \lambda \\ 0 &= \frac{\partial L}{\partial u} + \frac{\partial f}{\partial u} \lambda \\ \lambda(t_i^{(-)}) &= \lambda(t_i^{(+)}) \\ (\lambda + \nu)(T) &= \left. \frac{\partial \phi}{\partial y} \right|_T \end{aligned} \tag{16}$$

The boundary conditions are mixed with $y(0) = x_i$ and $\lambda(T)$.

These problems are notoriously hard to solve for nonlinear dynamical constraints due to the lack of knowledge (or educated guess) as to the initial value of the adjoint variable required for forward integration (or an approximate path needed to integrate backwards in time for the adjoint variable). In fact, trying to numerically integrate these equations with poorly guessed initial conditions for $\lambda(0)$ can "... produce 'wild' trajectories in the state space. These trajectories may be so wild that values of $x(t)$ and/or $\lambda(t)$ exceed the numerical range of the computer!" (Bryson and Ho, 1969).

C. Guiding Computation with Dynamical Systems Theory

It has recently been recognized that nonlinear trajectories provide a network of free trajectories that offer transport to widely separated regions in phase space (the 'IPS'). This provides an efficient means to sample from the initial ensemble (as we will demonstrate). In other words, we can map out the invariant manifolds in the phase space, and patch together trajectory segments.

Mapping out the IPS network holds the promise of maximizing the *reachable set* of configuration space at minimal *cost*. For space missions specifically, we want to maximize the volume of locations in the Solar System that are reachable within some time with a given fuel budget. It has recently been recognized that, when including the full non-linearity of the three and higher body dynamics, there are "free" trajectories (i.e. true solutions of the dynamics) that nevertheless result in transport to widely different regions of configuration space. Specifically, these trajectories tend to wind on and off various invariant manifolds in the phase space, and are suggestive of a new approach to computing solutions to optimal control, with the computation of trajectories guided by insight provided by dynamical systems theoretic views of the global structure of phase space. Provided we can meet the challenges of computing these trajectories and subsequently navigating them, *we are rewarded with longer mission duration and increased potential for science return*.

III. A Probabilistic Approach

A. Motivation

Independent of the algorithm used, we always have to supply an "initial guess" of the solution to the control problem which is to be improved with subsequent computation until some stopping criterion (we reach some target level of the optimality figure of merit or reach our maximum allowed time of computation). We refer to the initial guess solutions (we have an entire set for the set of initial and/or final states) as the *initial ensemble*. Assuming some distribution on the initial and final states, we have (even for deterministically supplied initial solutions) a *distribution of residuals* $P^{(0)}(\Gamma)$.

$$\begin{aligned} P^{(0)}(\Gamma) \quad = \quad & \int_0^\Gamma d\gamma \int d(x_i, x_f, T) \ p(x_i, x_f, T) \\ & \delta \left(\gamma - \Gamma(y^{(0)}(x_i, x_f) \right) \end{aligned} \tag{17}$$

where the initial solution approximation, denoted as a function of the boundary conditions is denoted $y^{(0)}(t; x_i, x_f)$. There is a *limiting distribution* given by finding the optimal solution for every set of randomly chosen BC's, $P^{(*)}(\Gamma)$ where

$$P^{(*)}(\Gamma) = \int_0^{\Gamma} d\gamma \int d(x_i, x_f, T)\ p(x_i, x_f, T) \\ \delta\left(\gamma - \Gamma(y^{(*)}(x_i, x_f)\right) \quad (18)$$

We therefore want an algorithm (deterministic or otherwise) which provably converges in distribution

$$P^{(n)}(\Gamma) \rightarrow_n P^{(*)}(\Gamma) \quad (19)$$

and even better if it does so quickly (i.e. in as few a number of "steps" as possible, i.e. with as little computational expense as possible).

It is difficult to prove convergence in distribution for deterministic systems, as for non-linear dynamical constraints we often encounter convergence to local minima, resulting in a distribution that is always *majorized* by the limiting measure (more concentrated about lower residuals, since each solution reaches a local minimum with residual greater than that of the global optimum for each set of BC's). The goal of convergence in distribution of the residual for random boundary conditions (or at the very least random initial conditions) leads us to consider a probabilistic approach to the entire problem itself. Our strategy is to 1) construct a sequence of probability measures directly which have the desired convergence properties as above, and 2) using sampling algorithms (including MCMC, particle filters, etc) to directly sample solutions from these probability densities. With this approach, we have an algorithmic framework for any dynamical system, which provably converges in the probabilistic sense above.

Finally, other sources of uncertainty can naturally be included in this framework, including perturbations to the dynamics, as well as the closed-loop control challenge of smoothing (quantifying uncertainty in the entire past trajectory, up to and including the current state) when supplied with state measurements with noise.

B. Components of the Probabilistic Approach

Roughly, the joint density for "everything" is

$$p(D_x, y_{1:n}, n, T, x_i, x_f) = p(D_x|y_{1:n}, x_i, x_f)p(y_{1:n}|x_i, x_f) \\ \times p(n|T)p(T|x_i, x_f)p(x_f, x_i)$$

where (D_x, D_u) are noisy measurements of the state and control inputs respectively, $y_{1:n}$ is a parametrized representation of a continuous time solution estimator (for example the state vector at the discrete times $\{t_{1:n}\}$, but other forms can be used as well), T is the total time, and (x_i, x_f) are the boundary conditions. The "inverse problems" of inferring trajectories given initial conditions, and/or control inputs given target final states and noisy state measurements along the way, are then viewed as solved by sampling from the associated conditional densities. What is interesting is that a very broad class of problems can be formulated with this specific mathematical

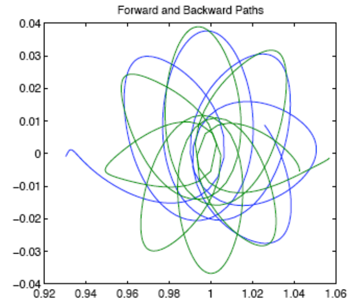

Fig. 1. Numerically integrated trajectories given by choosing random initial conditions about the CR3BP Lagrange (fixed) points $L1$ and $L2$, followed by integrating forward and backward in time resprectively.

goal - *sample from various marginalized conditionals from the above joint density*. We refer the interested reader to (Jewell, 2007, in preparation) for more details concerning the construction of the factors in the joint probability density above as well as proofs of convergence to sets of optimal solutions given randomly chosen boundary conditions.

We now turn to sampling from the initial ensemble, which is always used to initialize any algorithm of choice. As mentioned previously, our strategy is to use insight from dynamical systems theory to effeciently generate trajectories that have the rough qualitative character of optimal trajectories (possibly over a wide range of total times T).

IV. Sampling from the Initial Ensemble

A. Dynamical Systems Insight into Choice of BC's

Following (Gomez et al., 2002, Nonlinearity) we have very educated guesses for choosing initial conditions such that, when integrated forward and backward in time from the vicinity of the Lagrange points, pass near each other along a Poincare section at the plane $y_1 = 1.0$ From the dynamical systems theoretic viewpoint, these forward and backward integrated trajectories follow the unstable and stable (forward in time) manifolds, which we now know intersection in the plane $y_1 = 1.0$. Shown in figure 1 are such paths, which we will smoothly patch together to form a continuous path as an initial trial solution to the optimal control problem.

B. Approximate Trajectories

The intuition is that we can quickly generate reasonably good approximations to optimal solutions by creating a "library" of shorter length solution segments, for which dynamical systems theory and "global views of phase space" can provide guidance, followed by smoothly patching these segments together in order to satisfy postulated or assumed global constraints (such as paths which solve a control problem, agree with noisy measurements, etc.). Here we consider to manner in which we can patch together these solution segments - we phrase this as a control problem and provide a solution from a variational approach.

To do this note that for any path we have the required control inputs to follow that path given by

$$\dot{x}(t) = F \circ x(t) + u(t) \tag{20}$$

One way to find low residual paths is to consider paths built up by smoothly transitioning between local path segments of small residual (i.e. as generated by integrating forward and backward initial value problems). We can then simply define

$$\hat{u}(t) + \delta u(t) = F \circ x - F \circ y - [J \circ y](x(t) - y(t)) \tag{21}$$

where we have arbitrarily set the "control inputs" $\hat{u}(t)$ to be determined by some chosen path which satisfies the dynamics linearized about the reference paths

$$\dot{x}(\alpha_{1:n}, t) = \dot{y} + [J \circ y](x(\alpha_{1:n}, t) - y(t)) + \hat{u}(t) \tag{22}$$

subject to the boundary conditions. If we denote the residual of the reference path segments as ϵ, then we have the upper bound on the residual of our path

$$\|\dot{x} - F \circ x\| \leq \|\epsilon\| + \|\delta u\| \tag{23}$$

We are free to choose any path, but the intuition is that good choices are given by those that follow the reference paths but satisfy BC's while minimizing quadratic cost.

Specifically, the goal is to find paths, given a reference path (with jumps at discrete times), which minimizes the quadratic "defect" cost subject to the piecewise dynamical constraints

$$f(x, u, t) \quad = \quad \dot{y}_i + [J \circ y_i](x - y_i) + u \tag{24}$$

The overall functional to be minimized in order to solve this constrained optimization problem is

$$\begin{aligned} J \quad = \quad & \left(\int dt \ \|u\|^2 \right) + \nu \cdot x(T) + \sum_i \left(-\lambda \cdot x|_{t_i}^{t_{i+1}} \right. \\ & \left. + \int_{t_i}^{t_{i+1}} dt \ \lambda(t) \cdot f(x, u, t) + \dot{\lambda} \cdot x \right) \end{aligned} \tag{25}$$

Note that the dynamical constraints here are $\dot{\delta} = [J \circ y_i](t)\delta(t) + u(t)$, which gives the solution on each piecewise interval, for any driving force $u(t)$,

$$\begin{aligned} \delta(t \in I_j) \quad = \quad & e^{\int_{t_j}^{t} dt \ J \circ y_j(t)} \left(x(t_j) - y_j(t_j) + \right. \\ & \left. \int_{t_j}^{t} dt' \ e^{-\int_{t_j}^{t'} dt'' \ J \circ y_j(t'')} u(t') \right) \end{aligned} \tag{26}$$

Also note that for any exact solution to the dynamics $\dot{x} = F \circ x$, there is a driving force $u = F \circ x - \dot{y}_i - J \circ y_i(x - y)$. The idea is that for reference paths which are close to or exact solutions over the intervals I_j, that if we find paths that smoothly transition from one manifold to the other with the control input that minimizes $\|u\|$, then this will be a close to a true solution provided the jumps are small. This problem then allows us to patch together many short solution segments, which are "concatenated" together using insight from dynamical systems theory.

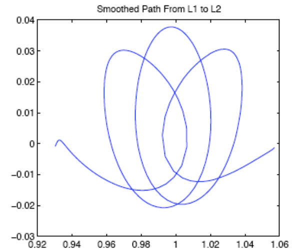

Fig. 2. The continuous path obtained by solving the optimal linear control problem given by matching specified boundary conditions about the two Lagrange points.

As discussed previously, calculus of variations gives the system of equations

$$\begin{aligned} \dot{x} &= \dot{y}_i + J \circ y_i(x - y_i) + u \\ \dot{\lambda} &= -[J \circ y_i]^T \cdot \lambda \\ u &= -\lambda^T \\ \nu &= \lambda(T) \end{aligned} \tag{27}$$

We know that in the absence of the endpoint boundary conditions (with $\nu = 0$), that $u = 0$, which is consistent with the above equations, and where forward in time the solution will track $y_1(t)$ up to the first jump, after which the fluctuation will evolve linearly during the next arc, and so on. Also, with no jump $x \equiv y$. Since we have a piecewise linear system, we might expect that the vector ν linearly depends on the initial and final state.

As noted previously, stability of the numerical solution of the above is a major concern. One approach to control this is essentially a stabilized march method (see Ascher et. al., 1995), with some minor modifications. The idea is to integrate $2n$ solutions to the optimal linear control equations forward and backward in time from the jump, with initial conditions given by variations in each of the standard basis. As we march the collection of fundamental solutions along, we monitor their linear independence (here measured by the condition number of the evolved covariance matrix). When a threshold is reached, we reset the initial conditions of the fundamental solutions to the original (therefore giving the original covariance matrix).

Using this procedure with the forward and backward paths, and a jump at the time of closest approach, we obtain the smoothed continuous path shown in figure 2.

V. Conclusions

Motivated by the potential increase in science return from space missions enabled by improvements in trajectory generation in the presence of uncertainty, we have formulated a probabilistic approach to trajectory generation which provably converges to entire sets of optimal solutions for random boundary conditions. Such a capability is particularly important in both the design stage (where typically a “sensitivity analysis” is important in order to validate the capability of the spacecraft to achieve the mission objectives even in the presence of launch vehicle error and other perturbations) and in flight, requiring precise control over uncertainty for tracking and navigation.

While optimal control problems are notoriously difficult to solve for nonlinear dynamical constraints (see for example Bryson and Ho, 1969), we meet this challenge with recent insight into the global phase space structure for dynamical systems relevant for space mission design (Howell et al. 1997, Serban et al. 2002, Gómez et al., 2004) as well as a probabilistic approach which provably converges to sets of (globally) optimal solutions. Specifically our strategy is to 1) use insight provided by dynamical systems theory to quickly sample from what we call the “initial ensemble” of solutions, 2) construct a sequence of probability measures directly which provably converge on sets of optimal solutions given randomly chosen boundary conditions, and 3) use sampling algorithms such as Markov Chain Monte Carlo or particle filters to directly sample solutions from these probability densities, with increasingly better solutions “learned” as computation progresses. In this paper we have demonstrated one approach to using dynamical systems theory to provide trajectory segments which can be smoothly joined (using optimal linear control) to quickly compute initial trajectories with the qualitative character of optimal solutions. Future work will use stochastic sampling algorithms to *improve the “cost” figure of merit in an algorithmic setting which provably converges*.

ACKNOWLEDGMENT

This work was supported under a grant from the NASA Applied Information Research Program.

REFERENCES

[1] K. Howell, B. Barden, and M. Lo, *Application of dynamical systems theory to trajectory design for a libration point mission*, J. Astronaut. Sci., 45, 1997, 161-178.

[2] K. Howell, B. Barden, R.S. Wislon, and M. Lo, *Trajectory design using a dynamical systems theory approach with application to Genesis*, AAS/AIAA Astrodynamics Specialist Conference (Sun Valley, Idaho), AAS paper 97-709.

[3] R. Serban, et. al., *Halo Orbit Mission Correction Maneuvers Using Optimal Control*, Automatica, 38, 2002, 571-583.

[4] G. Gómez, et. al., *Connecting Orbits and Invariant Manifolds in the Spatial Restricted Three-Body Problem*, Nonlinearity, 17, 2004, 1571-1606.

[5] A.E. Bryson, and Y. Ho, *Applied Optimal Control*, Blaisdell Publishing Company, Waltham, Massachusetts, 1969.

[6] U.M. Ascher, R.M. M. Mattheij, R.D. Russell, *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*, Classics in Applied Mathematics, vol. 13, SIAM, Philadelphia, 1995.